

Consistency in the face of change: an adaptive approach to physical layer cooperation

Ayan Sengupta^{*} Yahya H. Ezzeldin[†] Siddhartha Brahma[‡]
Christina Fragouli[†] Suhas Diggavi[†]

^{*}Stanford University, USA [†]UCLA, USA [‡]IBM Research - Almaden, USA

ABSTRACT

Most existing works on physical-layer (PHY) cooperation (beyond routing) focus on how to best use a given, static relay network—while wireless networks are anything but static. In this paper, we pose a different set of questions: given that we have multiple devices within range, which relay(s) do we use for PHY cooperation, to maintain a consistent target performance? How can we efficiently adapt, as network conditions change? And how important is it, in terms of performance, to adapt? Although adapting to the best path when routing is a well understood problem, how to do so over PHY cooperation networks is an open question. Our contributions are: (1) We demonstrate via theoretical evaluation, a diminishing returns trend as the number of deployed relays increases. (2) Using a simple algorithm based on network metrics, we efficiently select the sub-network to use at any given time to maintain a target reliability. (3) When streaming video from Netflix, we experimentally show (using measurements from a WARP radio testbed employing DIQIF relaying) that our adaptive PHY cooperation scheme provides a throughput gain of 2x over nonadaptive PHY schemes, and a gain of 6x over genie-aided IP-level adaptive routing.

1. INTRODUCTION

Physical layer (PHY) cooperation can help relieve the bandwidth crunch that remains a core threat to mobile user experience today. A number of theoretical works have established the significant benefits enabled by PHY cooperation; system implementations have also validated that it is feasible for current devices to support PHY cooperation [8, 5, 6]. However, a crucial missing piece remains: how does one implement PHY adaptation over cooperative networks? Most works beyond routing have focused on optimally operating a given static relay, while today, in our homes, at work and in public places, we can have changing network topologies and multiple wireless (*e.g.*, WiFi-enabled) devices within range to assist a source-destination pair.

Consider for instance the following scenario: Alice is streaming Netflix to her tablet using WiFi, and takes the tablet with her as she moves around her house¹. Her

¹Our application focus is on indoor scenarios, where we do

not expect “fast-fading”, as can occur in (LTE-based) cellular networks for objects moving at high velocities.

house has a WiFi router wired to the Internet that acts as source, as well as three other wireless devices that can act as WiFi relays, and can implement PHY cooperation with the source. As her network conditions change, which (if any) of the three potential relays should assist her, so that she can experience a consistently good video quality? This is a well studied problem in routing, where one selects which relays to route packets through; in this paper we focus on PHY cooperation networks, where there is (potentially) a larger gain over adaptive routing, yet practical algorithms for exploiting the same are less well understood.

We are interested in the following questions: Should we select one or more relays to assist us via PHY cooperation? What are the best relay(s) to use? Is it easy to identify them? How can we efficiently adapt our choice as Alice moves? How important is it, in terms of performance, to adapt? And what are the benefits of adaptive PHY cooperation over adaptive routing? Answering such questions is a necessary step in bringing PHY cooperation closer to practical systems.

We restrict our attention to using either none, one or two relays, and at most two cooperating transmitters at a time. We make this choice for two reasons. First, we do not expect significant benefits from simultaneously using three or more relays; this is supported by numerical evaluations of outage performance (in Section 3.3) that indicate a “diminishing returns” trend for larger number of relays—the more relays we use, the less we gain per relay. Second, the complexity of the system increases exponentially with the number of simultaneously transmitting nodes, as we discuss in Section 3.4.

We build upon our PHY cooperation scheme DIQIF [5] and customize it for two relay cooperation. We explore the benefits that two relays can enable in practice.

Next, we introduce SPA, an algorithm that selects whether to use a subset of one or two relays, and exactly which subset, from among a set of available relays, by performing Frame Error Rate (FER) learning. Our contribution is a design that carefully balances the train-

not expect “fast-fading”, as can occur in (LTE-based) cellular networks for objects moving at high velocities.

ing overhead with the estimation performance of SPA, and gracefully adapts to channel changes. Inspired from machine-learning techniques, SPA learns the best subset while sending the actual data we want to transmit (as opposed to dummy training data), invokes an early reject criterion while learning to quickly weed out relays that are not likely to lead to good performance, and uses memory to adapt to changing network conditions. In addition, SPA is completely agnostic to the underlying PHY relaying scheme, and can be used in conjunction with other strategies like Decode-Forward, Amplify-Forward etc., if desired.

Our next contribution is a proof-of-concept deployment on a WARP software radio testbed. We find that DIQIF outperforms alternative schemes for two-relay cooperation; that performance is significantly affected by the choice of cooperating relay(s); that SPA can reliably select an operational mode that enables consistently low FERs in the face of changing network conditions; and that the benefits from adaptation in SPA can enable Netflix streaming at higher (and consistent) rates, promising a better user experience. Interestingly, when pitted against protocols without PHY cooperation, SPA significantly outperforms ($\geq 6\times$ Netflix throughput gain) even a genie-aided adaptive routing strategy with knowledge of future network conditions.

In what follows, Section 2 reviews related work; Section 3 presents the relaying scheme and theoretical evidence through numerical evaluation; Section 4 introduces SPA, our algorithm for relay selection; Section 5 describes system implementation; Section 6 presents experimental evaluations; Section 7 tests video streaming performance, and Section 8 concludes the paper.

2. RELATED WORK

Relay Selection Algorithms: Theoretical algorithms (of both distributed as well as centralized nature) for relay selection are studied in [10, 23, 22, 13, 17] and the references therein. In [10, 23, 13, 22] orthogonalized single-carrier transmissions, as well as Amplify-Forward (AF) and/or Decode-Forward (DF) relaying strategies are considered and algorithms are proposed that use an objective function based on channel state information (CSI). In [17] a version of the Partial DF [24] relaying scheme is studied and algorithms based on CSI as well as distance are proposed. However, the objective functions for relay selection in [17] are based on high-SNR approximations. These theoretical works do not offer testbed evaluations of their algorithms in real-world scenarios. Moreover, the schemes considered—AF and DF, have also been shown to perform suboptimally in theory [2] as well as practice [8]. Additionally, and differently from our approach, [17, 13] assume accurate forward CSI at the transmitter which can be difficult to acquire in wireless systems and also represent an overhead bot-

tleneck in multicarrier (OFDM-based) systems, while introducing errors in SNR-limited regimes. Multiuser scheduling algorithms for exploiting cooperation benefits from OFDM-based DF relay networks are studied in [7]. Hop-by-hop best relay selection techniques for exploiting diversity in routing-based multihop wireless networks are explored in [12].

Network Information Theory: Our work is inspired by the information-theoretic result on wireless network simplification [19], which shows that using subsets of relays can achieve a constant fraction of the capacity; [19] provides capacity results for static Gaussian channels, while we focus on adaptation over time-varying configurations and practically implementable relay selection protocols. [2] presents the QMF scheme that approximately achieves the network capacity and which is the foundation for the DIQIF scheme [8] used in this paper.

PHY Cooperation Testbeds: For concurrent sender-receiver pairs, Analog Network Coding and interference alignment based system implementations were proposed in [14] and [9], respectively. The first testbed implementation of QMF over single-relay systems, as well as that of coded non-orthogonal AF and DF relaying systems were presented in [8], demonstrating significant benefits of QMF over AF and DF. [5] extended the work of [8] to include opportunistic decoding/quantization at the relay and hybrid decoding at the destination. Other WARP [15] based implementations of PHY cooperation networks include the work of [18], where uncoded AF and DF relaying were used. MAC-PHY layer protocols for optimally triggering cooperative transmissions in WiFi networks were studied in [4, 21, 3].

3. COOPERATIVE TRANSMISSION SCHEME

We here first describe the network operation and give an overview of DIQIF [5]. Our new contributions are that we customize DIQIF to 2 relay networks, and present theoretical evidence to support our design choices.

3.1 Network Operation

Modes. Consider a source S that wants to communicate with a destination D . We have a direct link connecting S to D , as well as N relays available to help the S – D pair. Our goal is to adaptively select which 1 or 2 relays, among the N , will assist the S – D communication if needed, at each point in time. Clearly, we have $N + \binom{N}{2}$ choices; we term these choices “modes”. We have N “1-relay modes” where the source cooperates with a specific relay i , and $\binom{N}{2}$ “2-relay modes” where the source cooperates with two specific relays i and j .

Two phase operation. In both types of modes, the source first attempts to communicate with the destination; if




Figure 1: Two-Stage Network Operation

it fails, a cooperative transmission takes place:

- *1-relay mode*. In Phase 1 the source transmits; the destination and the one relay in the mode listen to the transmission. If the destination cannot decode, in Phase 2, the source and the relay cooperatively transmit; the destination receives the superposition of the source and the relay signals.
- *2-relay mode*. In Phase 1 the source S transmits; the destination D and the two relays listen to the transmission. If the destination cannot decode, in Phase 2, the two relays cooperatively transmit; D receives the superposition of the two relay signals.

Fig. 1 depicts, for $N = 2$, the three possible modes and the two-phase operation, where one or both relays, respectively, assist the S–D communication. Since we restrict our attention to two cooperative transmitters at any time, we do not consider the case where the source would transmit simultaneously with the two relays; we discuss why in Section 3.3 (P4) and Section 3.4.

3.2 Relaying Strategy

We implement DIQIF relaying over the PHY procedures of WiFi, such as OFDM modulation. The source transmission is always a coded information stream using standard (for WiFi) LDPC codes.

DIQIF [5] for a 1-relay mode operates as follows. The relay first attempts to decode its received signal. If it succeeds, then the relay has access to the clean information bits. It re-encodes the information bits with the same code as the source, and interleaves the coded bits. If it does not succeed in decoding, the relay quantizes the elements of its received signal to their closest constellation points, and interleaves the resulting bits. The relay and the source then synchronously transmit the coded sequence (after appropriate OFDM modulation) in Phase II (if Phase I fails), effectively creating a distributed space-frequency code.

To extend this scheme to two relays, we apply the same operation (if possible decode, otherwise quantize) at each relay, and have the two relays synchronously transmit in Phase II. However, we ensure that each relay uses a *different* interleaver, as we found that this is crucial in achieving diversity, and thus a good error performance. For decoding at the destination, we

designed a belief-propagation joint decoder, that takes into account the structure of the (two-relay) network (including quantizers at the relays, and superposition at the destination), and jointly decodes from the (different) relay transmissions of the source message.

3.3 Theoretical evidence

We here provide evidence, through theoretical modeling and simulations, to support the following points:

- P1*. There can be a significant performance variability depending on which relay(s) we select.
- P2*. As long as we select a few of the relays carefully, we can already extract a significant portion of the cooperation benefits that the full network offers.
- P3*. As the number of relays increase, we get diminishing returns in terms of additional benefits.
- P4*. In a 2-relay network, having both relays and the source simultaneously transmit does not offer significant benefits over the best 2-transmitter mode.

We emphasize that this evidence is only indicative of the trends we expect in our testbed: analyzing the performance of cooperative networks that utilize the signal constellations, practical relaying schemes and codes as in our implementation, is not theoretically tractable; therefore we resort to the closest tractable theoretical models, that could still enable us to gain some intuition.

We here summarize the steps and assumptions of the theoretical analysis; we provide more details in Appendix A. The relevant information-theoretic metric over Gaussian fading channels is outage probability: the probability that the channel realizations do not support a target rate R . To calculate the rate R , we use the approximate capacity expressions in [1, 20] for Quantize-Map-Forward (QMF), an information theoretic scheme, that has ideologically inspired (in terms of operation) DIQIF [5]. QMF uses Gaussian codebooks for transmission and Gaussian quantizers to quantize each received vector at the noise level, followed by random mapping and retransmission. We assume infinite complexity encoding and decoding at the network nodes. We write expressions for outage probability using channel statistics, and formulate an optimization problem that selects the best subset of relays (of size 1, 2, 3 etc.) to use over a given N -relay network.

We numerically solved the optimization problem over a range of configurations, assuming Rayleigh fading; we provide indicative plots in Fig. 2, over a network with 10 relays, where either none, or k ($= 1, 2, 3, 4$), or all relays are chosen. Fig. 2(a) and 2(b) consider a clustered topology where some of the relays are in close range of the S–D pair while the rest are further spread out. Fig. 2(c) corresponds to a symmetric setting, where all network channels have the same fading statistics.

To support *P1*, Fig. 2(a) shows that the relays we select can make a significant difference, by comparing the




Figure 2: Outage performance over Rayleigh-faded relay networks with increasing number of relays selected out of maximum of 10 relays. The x -axis denotes *per node* reference SNR.

performance of the optimal 1 and 2 relays (i.e., relays that leads to the best performance) to other suboptimal choices. To support $P2$ and $P3$, Fig. 2(b) and 2(c) show, over a clustered and a symmetric configuration, the fact that the returns from increasing the number of relays diminish fast. For instance, in Fig. 2(b), at an FER of 10^{-2} , we see that going from 0 relays to 2 provides a 13 dB gain in performance, whereas the benefit in going from 2 to 10 relays is 4 dB in total.

To support $P4$, Fig. 3 compares the performance when all three nodes (S , R_1 and R_2) in a 2-relay network simultaneously transmit, vs when any two of the nodes do so. We see that, while a particular 2-transmitter mode can perform worse than using all 3-transmitters, the best incurs minimal performance loss for SNR ranges of interest. Intuitively, if e.g., the S - D link is strong, then having S and the “strongest” relay transmit can perform very close to all three nodes transmitting; similarly, if the S - D link is weak, the 2-relay mode approaches the optimum performance. Hence, if we can select the best 2-transmitter mode, we can achieve comparable performance to that with all 3 transmitters active.

3.4 Discussion

Why DIQIF. We selected DIQIF for the relay operation for two reasons. First, it was shown to be the best PHY strategy in 1-relay networks [5]; in this paper, we extend DIQIF to support 2-relay modes and experimentally verify (in Section 5) similar trends. Second, with DIQIF, each relay performs the same operations independently of whether it operates in a 1-relay or a 2-relay mode, and independently of channel statistics. This is an attractive attribute when we need to adapt the relays as the network changes. We note that, although we selected a specific relaying scheme that leads to good performance, the adaptation algorithm we describe next in Section 4 could be applied over any relaying scheme, such as Amplify-Forward, Decode-Forward, etc.

Using more than 2 transmitter cooperation. As Fig. 2 shows, the outage performance could be improved by adding more than two relays. However, we believe that

the overhead for more than 2-transmitter cooperation would outweigh the benefits, as we argue next.

- **Complexity of Decoding:** Unlike setups such as MIMO, where we can perform beamforming and successively decode multiple superimposed signals, we have a single receiver antenna and a superposition of different signals that need to be simultaneously processed to extract the source message. As a result, the complexity of the decoder increases exponentially with the number of superimposed signals; although we used a customized belief propagation decoder that offers the state of the art in low-complexity decoding, more than two superposed signals led to prohibitive processing in our decoding.

- **Overhead:** For PHY cooperation transmission, preamble information (training sequences and SNR estimates) needs to be transmitted from each of the k -transmitters to the receiver orthogonally over time. Thus, for a fixed payload size, the preamble overhead to payload scales linearly with the number of transmitters used².

- **Power Efficiency:** Having a small number of transmitting nodes is also attractive from the standpoints of interference, total power efficiency, and even total indoor emissions (health regulations). Fig. 4 plots the performance of best k -relay subnetworks, where the x -axis is normalized to the total power used for transmissions by all nodes in the network. We see that a total power budget is often better utilized if it is distributed among a smaller subset of well-chosen relays, as opposed to distributing it across all relays.

- **Diminishing Returns:** Fig 2 (without factoring in overheads) shows that the progressive benefits in going from 2 to 3 relays are 50 % of that in going from 1 to 2 relays, and 17% of that in going from 0 to 2 relays.

4. SPA: SELECTING RELAYS

Given a network with N relays, SPA selects which mode to activate, i.e., which one or two relays will best serve to assist the S - D communication. We assume that S transmits at a fixed rate, e.g. to support video streaming. During selection, we use a fixed MCS (determined by the source, and also used by the relay), and employ

²In Section 5.2, we see that the overhead ratio in our setup is approximately $0.07k$ for k -transmitters.




Figure 3: Performance of 2-transmitter modes in Rayleigh-faded 2-relay networks.




Figure 4: Power efficiency in a 10 relay network.

end-to-end FER as our performance metric. Since the discovered ordering (which mode is better) extends to any fixed MCS, with rate adaptation (on top of mode selection), the best mode still achieves a better FER.

4.1 Main idea

SPA starts by identifying the mode that achieves the lowest FER among all $N + \binom{N}{2}$ modes. Subsequently, SPA is triggered whenever the experienced FER of the selected mode exceeds a predefined threshold. The main design challenge is efficiency: even for a moderate N , learning across all modes at the destination may incur significant overhead - both in terms of rate and in terms of the “mode switches” required to test the modes and converge. We make SPA efficient in three ways:

- We use information-carrying frames for training, i.e., we send information messages via each of the modes and use the FER of the received information messages³ to decide on the winner mode during learning. We refer to these as training frames in the following, yet we emphasize that there is *no rate overhead*; only the performance penalty of not always utilizing the best mode.
- Instead of sending the same number of training frames for each mode, we progressively weed out modes with high FER and thus allocate more training resources to the modes that are likely to be selected. Thus we converge faster to the mode with the lowest FER.
- SPA keeps memory—in the form of ranking of modes after each learning phase. The intuition is that if the configuration changes, the “now best” mode is likely to be found among the previously stronger modes. Accordingly, every time we need to adapt, we use the ranking of the modes in the previous learning phase to train over

³We estimate the FER by decoding and checking the CRC.

a smaller *partition* (subset) of modes. This reduces the number of “mode switches” required, and as we show in Section 6 (Fig. 9(c)), suffers from negligible FER-penalties. Also, it is this facet of the algorithm that gives it its name: **Sort**, **Partition** and **Adapt**.

4.2 Algorithm description

SPA is described in pseudo-code in the following page. For clarity, we separately describe the function LEARN that searches for the best mode starting from a set of candidates; SPA invokes LEARN to adapt as needed, and provides to LEARN the starting set of mode candidates. SPA is inspired from a popular machine learning approach of *learning from the best expert* [11], appropriately customized for our setup. Two main differences are that here we are looking to *converge* to the best expert’s FER, where an expert corresponds to a mode; and that we use memory to adapt to network changes.

LEARN Find the best mode

Input: Set of r modes $\{m_1, \dots, m_r\}$.

Output: Ranked list of r modes.

$C_0 \leftarrow \{m_1, \dots, m_r\}, \forall i \in C_0, w_i \leftarrow 1/r, j \leftarrow 0, L' \leftarrow \{\}$.

while ($|C_j| > 1$ and $j < B$) **do**

Run the modes in C_j for l frames.

Compute the empirical FER for mode i as \hat{f}_{ij} .

$P \leftarrow 0, tot \leftarrow 0, C_{j+1} \leftarrow \emptyset$

for $i \in C_j$ **do**

$w_i \leftarrow w_i e^{-\eta \hat{f}_{ij}}$

$P \leftarrow P + (1 - (1 - \alpha)^{\hat{f}_{ij}}) w_i$

end for

for $i \in C_j$ **do**

$w_i \leftarrow (1 - \alpha)^{\hat{f}_{ij}} w_i + \frac{1}{n-1} (P - (1 - (1 - \alpha)^{\hat{f}_{ij}}) w_i)$

$tot \leftarrow tot + w_i$

end for

for $i \in C_j$ in ascending order of w_i **do**

$w_i \leftarrow w_i / tot$

if $w_i > \epsilon$ **then**

$C_{j+1} \leftarrow C_{j+1} \cup \{i\}$

else

$L' \leftarrow \text{append}(L', \{i\})$

end if

$j \leftarrow j + 1$

end for

end while

Sort C_j in ascending order of weights.

Append C_j to L' and output $\text{reverse}(L')$.

Learning in batches and weight update. LEARN

starts from a set of admissible modes. There is a weight w_i for each mode i . Initially all weights are equal, with their sum normalized to 1. Learning takes place in *batches* of frames. In each batch, l frames per admissible mode are sent for training (we emphasize that these frames encode and convey information messages to the destination; they are not dumb training frames). After each batch, the weights of the admissible modes are updated. A mode with high FER is penalized exponentially in terms of its weight; the learning rate η decides the rate of penalization. The shifting parameter α enables robustness against sudden changes in FER.

Early Reject of Bad Modes. After each batch is processed, the modes with w_i 's less than a threshold ϵ are rejected. This ensures that (i) bad modes are not retained throughout the learning phase, and (ii) when a clear winner exists, faster convergence to the winner.

Mode Selection and Termination. Learning can at most continue for B batches, after which a hard decision is taken by selecting the mode with the highest weight. The algorithm may converge before B batches if *exactly one* mode remains in the learning phase. We set the value of B online during training, to a value for which the algorithm converges; this incurs no extra overhead.

Adapting to configuration changes. LEARN is used by SPA to adaptively select the best mode of operation based on a triggering mechanism. LEARN is invoked whenever the “windowed FER” (i.e. FER in the last w frames) of the currently used mode exceeds a pre-determined threshold (ζ).

Learning over ranked subsets of modes. Whenever we invoke LEARN, we use the ranking of modes in the previous cycle to select a subset of “high-ranked” modes over which we run our algorithm. The subset size (r) is an algorithm parameter, providing a trade-off between accuracy (better for larger r) and switching overhead (better for smaller r). We maintain a ranked list L of modes, approximately sorted in ascending order of FER. Whenever the FER of the last w frames, which itself is computed at intervals of Δw frames, drops below ζ , we do either of two things: if the new trigger happened very quickly after the previous one (this is controlled by a parameter s), the top r modes in L are pushed to the end of the list; otherwise we only push the top mode to the end of L . In both cases, LEARN is called with the r modes on top of L , and their ranking is updated using the output of LEARN. After every invocation of LEARN, the top mode in L is used to operate the network. Table 1 collects the parameters used.

4.3 Discussion

Why use FER-driven learning? We consider part of our contributions, that using FER-driven learning in SPA is an educated choice: we put effort in exploring and ruling out alternative methods. We can think of

SPA Adaptive learning

```

Maintain  $L$  as list of modes in  $M$ .
while End of Time do
   $i \leftarrow 0$ .
  Run network using mode  $L[1]$  for  $w$  frames.
  while True do
    if Empirical FER in last  $w$  frames  $\geq \zeta$  then
      Break.
    end if
    Run network for  $\Delta w$  frames using mode  $L[1]$ .
     $i \leftarrow i + 1$ .
  end while
  if  $i \leq s$  then
     $L \leftarrow \text{append}(L[r + 1 \dots |M|], L[1 \dots r])$ .
  else
     $L \leftarrow \text{append}(L[2 \dots |M|], L[1])$ .
  end if
  Call LEARN with  $L[1 \dots r]$  and update top  $r$  entries of  $L$ 
  with its output.
end while

```

M	Set of modes where $ M = N + \binom{N}{2}$.
ζ	FER threshold
r	Memory size
w	Size of window to calculate empirical FER
Δw	Number of additional frames
s	Minimum number of windows
l	Learning batch size
η	Learning rate
α	Learning shift parameter
ϵ	Learning threshold
B	Maximum number of learning batches

Table 1: Parameters used in SPA and LEARN

three basic approaches for selecting relays: distance-based, channel-based and FER-based. For PHY cooperation, distance-based selection turns out to be too simplistic: it does not give consistent performance over wireless, due to effects such as interference, obstructions, multipath fading, etc. The second alternative is to learn the channels between all the network nodes, and use that knowledge to select relays, for instance, by evaluating theoretical rate expressions over the subnetworks. This approach, that works well for packet routing over relays, is surprisingly not reliable when selecting subnetworks for PHY cooperation. We ruled out this alternative because, from the channel knowledge, we can only calculate theoretically *approximate* expressions of the relay network performance; we found that the ordering (which subnetwork is better) indicated by these approximate expressions is not consistent with the ordering achieved by the actual scheme in practice.

For example, in Fig. 5, we see that over this static configuration, mode $R2R3$ is theoretically predicted (using channel estimates from the implementation) to be much better than $SR1$ and could support a 0.5 bps/Hz higher rate. However, looking at FER, we see that in fact it for $R2R3$, it is approximately 2x that of mode $SR1$. The inconsistency in the example is because, in contrast to point-to-point links, there does not exist exact information-theoretic capacity expressions for




Figure 5: FERs vs channel-based predictions.

PHY cooperation networks; only bounds on capacity are available, and the gaps are in general channel and topology (1 vs 2 relay(s)) dependent. Explicit theoretical characterizations of the QMF rate that take into account the quantization, mapping and channel estimation errors that occur in implementations such as ours, have so far remained elusive. Finally, this approach (similarly used in [13, 22, 17]) comes with the overhead of collecting (and sharing) accurate channel information of all channels, which is especially high in OFDM-based systems like WiFi, where we need channel information per subcarrier, per link.

How does SPA scale? SPA can be applied over arbitrarily large networks, since its complexity is related to the number of modes invoked in LEARN, as opposed to the total number of modes in the network, which is $N + \binom{N}{2}$ for an N -relay network.⁴ We found in our experiments that we do not need to train over more than 3 or 4 different modes at a time. For instance, in Section 6.3 (see Fig. 9(c)) we show that over 3 relay networks, a subset size of 3 modes (as opposed to the total of 6 modes) at each learning epoch allows to extract excellent FER performance while also minimizing the number of mode-switches necessary, achieving faster convergence and less feedback overhead.

Why not brute force search? Even with a small number of relays, SPA can significantly outperform exhaustive search (0.5x FER), as we show in Section 6.

5. SYSTEM IMPLEMENTATION

5.1 Physical Layer

We follow the PHY procedures of WiFi (IEEE 802.11). Each transmitted frame consists of a preamble and a payload, as illustrated in Fig. 6.

Preamble. The preamble consists of TAGC, TSYNC, and TCHE OFDM symbols structured as explained in Fig. 6. The TCHE OFDM symbol is additionally used for CFO estimation. In Phase 1, only the source transmits, using the structure in Fig. 6. In Phase 2, TAGC is sent by the two transmitters simultaneously, with a

⁴Having said that, we do not expect in an indoor environment (our application focus) an arbitrarily large number of available relays within range.




Figure 6: Time diagram.

cyclic shift between them to avoid accidental nulling. TSYNC and TCHE are orthogonalized in time. For k cooperating transmitters, the preamble therefore consists of $3k + 1$ OFDM symbols.

Payload. The payload consists of 42 OFDM symbols, i.e., data and pilot subcarriers as in 802.11. In Phase 1, for all schemes, we transmit the payload corresponding to an OFDM-based single Tx single Rx antenna system. In Phase 2, the source (if selected to transmit) retransmits the same payload as in Phase 1, and the relay(s) transmit processed signals. The preamble-to-payload ratio is therefore $0.07k$ for k cooperating transmitters.

Synchronization. For Phase 2 transmissions, carrier and timing synchronization is performed through a wired connection between nodes—the same approach as presented in [8]. Yet, we would like to mention that work on distributed transmissions has shown the viability of achieving accurate timing and carrier synchronization in a distributed manner [18, 21, 3]; these are enabled by implementing a large part of the mechanisms in real time in the FPGA to achieve fast turnaround times. We did not incorporate this into our testbed, where we focused on proof-of-concept experimentations.

5.2 Testbed

We used the WARP SDR hardware to implement the source, relays and destination nodes and the WARPLab framework to interact with the hardware via a host PC running MATLAB. The samples to be transmitted by a node were generated in MATLAB and downloaded to the transmit buffer of the corresponding node. The host PC triggered a real-time over-the-air transmission and reception by the nodes. The samples received at a node were read by the PC and processed in MATLAB.

6. EXPERIMENTAL EVALUATION

6.1 Performance over 2-relay networks

We created 4 scenarios, each having a source S, a destination D and two relays R1 and R2. We report the Received Signal Strength Indicator (RSSI) values for each scenario in Fig. 7(a); RSSI variations are due to distance, multipath, and transmit power adjustments. For each setting, we ran experiments for at least 1500 coded frames. We used 16-QAM constellations with a coding rate of $5/6$. In scenarios 1 and 2, we have R1




Figure 7: Results for 2-relay networks.

close to the source and R2 close to the destination; the difference is that in scenario 1 the S-D link is very weak. In scenario 3 the direct S-D link is again weak while the rest of the links are of approximately the same strength. In scenario 4 we have a strong S-R2-D path.

We compare: (i) *Decode-Interleave-Forward (DIF)*: Decode at the relay if possible and transmit the bit-level interleaved signal. If relay decoding is not possible, the relay remains silent. (ii) *Decode-Interleave-Quantize-Interleave-Forward (DIQIF)*: DIF if relay decoding succeeds; quantize to constellation, interleave and forward (QIF) otherwise. (iii) *Direct Transmission (DT)*, where in Phase 2, the source repeats the Phase 1 signal.

(1) DIQIF performance in 2-relay modes. Fig. 7(b) looks exclusively at the 2-relay mode, and compares the performance of DIQIF with that of DIF, in terms of FER. In keeping with the trends from 1-relay networks in [5], DIQIF exhibits consistently better performance than DIF. This is because, even when the relay cannot decode, with physical layer cooperation (and DIQIF) it can still forward useful (quantized) information that can help the receiver decode; in contrast, schemes like DIF, or simple routing along paths, require that the relay itself decodes, creating an unnecessary bottleneck.

(2) 0, 1 or 2 relays? Which ones? Fig. 7(c) shows the FERs if we operate: (i) using only the direct S-D link, (ii) constantly in 1-relay mode with relay R1, (iii) constantly in 1-relay mode with relay R2, and (iv) constantly in 2-relay mode with both R1 and R2. We report these for each of the 4 scenarios. We find:

- *Cooperation offers benefits over direct transmission.* As shown in Section 3, the maximum gains are obtained during the first steps of cooperation, i.e., including 1 or 2 relays, as opposed to none. Fig. 7(c) shows in some cases a *greater than two orders of magnitude benefit* in FER of the best cooperative mode over using DT.
- *We cannot select one mode to operate across all configurations.* This is because: (i) There is no mode that is universally better or is universally worse. In the spectrum of configurations (not all reported here), we found all modes to dominate in at least some particular scenario. (ii) The 1-relay mode can perform better than the 2-relay mode; this was the case in scenarios 2 and 4. This could happen when the S-D link is strong.




Figure 8: Adaptation in a time-varying network.

6.2 Adaptation provides benefits

Fig. 8 demonstrates the benefits of adaptive mode selection in a 3-relay network, when the topology of the network undergoes multiple changes over time. In this setup, we had 4 physical topology changes, and each topology was held constant for 172 frames. The relaying operation used was DIQIF, coding rate was 5/6 and modulation was 16-QAM. The algorithmic parameters for SPA were: $\zeta = 0.1$, $r = 3$, $w = 40$, $\Delta w = 1$, $s = 3$, $l = 1$, $\eta = 3$, $\alpha = 0.4$, $\epsilon = 0.05$, $B = 50$. We see that no mode without adaptation can on its own achieve the FER performance of SPA. This is because, which mode is best depends on the topology of the network, which changes over time – and SPA adapts to these changes.

6.3 Evaluating the design of SPA

To provide meaningful insights into the ingredients in SPA, we adopt the following evaluation methodology. We perform over-the-air experiments over 3-relay networks in 10 different topologies and run 860 frames in every topology for *all* 6 modes of operation in a time-interleaved fashion. From the experimental data (FERs of the individual modes) for these topologies, we create “samples” of time-varying 3-relay topologies, where each sample comprises 4 topology transitions and with 172 frames between transitions. The topology at each transition in a sample is chosen uniformly at random, with repetition, and the specific frames are chosen randomly from the (larger) dataset for that topology. We term a collection of such samples of time varying 3-relay networks as an *ensemble*. We run each of the algorithms over the samples in such an ensemble, and present *ensemble average* results for FER and switching performance to cover as wide a gamut of time-varying 3-relay networks as possible. In all the experiments, the relaying operation used was DIQIF. The coding rate, mod-




Figure 9: Ensemble average performance of different variants of adaptive mode-selection algorithms.




Figure 10: Comparison with Baseline Schemes

ulation, and the parameters for *SPA* were the same as in Section 6.2, unless otherwise stated.

We compared against the following set of algorithms.

Baseline Algorithms. We implement the following:

- *DT*: Non-cooperative strategy, where in Phase II, the Source simply retransmits to the Destination.
- *BRUTE* (Exhaustive Search): At every trigger, the empirical FER of every mode is measured, and a decision on the best mode made after observing a set number of frames from each. No machine-learning is used.
- *RandPick*: At every trigger (batchwise FER above a threshold), a *random* mode is selected for operation.
- *PWR2*: At every trigger, two modes are picked at random. The empirical FER of each is measured and the better is selected (Power of Two Choices Algorithm).

Adaptive Algorithms. The following make use of the machine learning techniques outlined in Section 4:

- *NRNM*: At every trigger, learning takes place with (i) No Early Reject (NR), and (ii) No Memory (NM), i.e., learning happens across *all possible* modes.
- *WRNM*: At every trigger, learning takes place (i) With Early Reject (WR), and (ii) No Memory (NM)
- *SPA*: At every trigger, learning takes place (i) With Early Reject (WR), and (ii) With Sorted Memory (SM), i.e., over a *ranked subset* of modes.

Comparison with baseline schemes. Fig. 10 demonstrates the ensemble average FER benefits of *SPA* over certain baseline strategies in 3-relay networks with multiple topology changes over time. In our topologies, the direct link was very weak, and it is hence another clear demonstration of the need for using physical layer cooperation to ensure connectivity. As we can see, *SPA* provides an FER that is 50% of the FER achieved by the best baseline strategy, which in our case, is *PWR2*. We would expect this gain to be even higher in scenarios

(‘samples’ from our ensemble) where there is a greater variability in the FER across the modes, and choosing bad mode(s) can significantly affect performance. The ensemble averaging technique used to present the results here, inherently “smooths over” the sample variability of *PWR2* or *Randpick*’s performance.

Early rejection and learning over subsets. Fig. 9 depicts ensemble average FER and switching performance for different variants of our learning-based algorithms for adaptive mode switching. Fig. 9(a) demonstrates that invoking the early-reject criterion and not allocating equal learning resources to all modes, provides a significant benefit in performance: the NRNM scheme suffers a factor of 2.5 penalty in average FER over WRNM and SPA, which employ the early reject criterion. This is because the early-reject allows a faster convergence to the requisite mode, while not sending too many frames via the obviously bad modes during the learning process. Fig. 9(a) also demonstrates that SPA, which learn over subsets (of size 3) as opposed to WRNM (which learns across all modes) only suffers a very nominal penalty (less than 0.4%) in FER performance. This is because convergence speed is much faster over smaller subsets, and even if a selected mode does not satisfy the FER requirements, the trigger mechanism allows the algorithm to quickly learn from the other partition of modes to converge to the best. In addition, a significant benefit of learning over subsets is observed in the average number of mode switches that are required in the course of running an algorithm over a time-varying 3-relay network. From Fig. 9(b), we can see that the mode-switching overhead reduces by a factor of ≈ 3 when learning takes place over subsets, again owing to faster algorithmic convergence over smaller sets. Reducing switching overhead is key to network operation, as this means less feedback requirements, less need of (re)-synchronization across terminals, and better network resource utilization in the sense of allocating free relays to other source-destination pairs.

Effect of subset size (memory) on learning. Fig. 9(c) depicts ensemble average FER versus switching performance with varying memory (or subset) sizes for learning. For 3-relay networks, we have at most 6 modes




Figure 11: Experiment setup for video streaming.

of operation, and a memory size of r (ranging from 1 to 6) in SPA for learning across the top r ranked modes, every time there is a trigger. We observe in Fig. 9(c) that average FER performance starts to improve with increasing memory size before saturating. The maximum gains are observed when increasing the memory size from 1 to 2 (about 44%). Switching overhead, on the other hand, increases monotonically and significantly with memory size. While a smaller memory size warrants less mode switches during learning, it also suffers from the penalty of sending too many frames through a badly chosen mode before the algorithm can detect a mode that meets the FER requirements. The tradeoff here is interesting as it seems to suggest that as long as we are operating above a certain memory size (3 in the case of our networks), the differences in FER performance are insignificant, while we gain a lot in terms of switching overhead with using a smaller acceptable memory size (factor of ≈ 2 from size 3 to size 6).

7. AN APPLICATION PERSPECTIVE: VIDEO STREAMING

In this section, we evaluate how the benefits of adaptive PHY cooperation translate to video streaming performance and also compare it with adaptive routing.

7.1 Experiment Setup

We implement the setup in Fig. 11: Alice runs a Netflix web browser client on her device (in our experiments, a Linux PC) using an 802.11g wireless router with bit-rate of 6 Mbps⁵; we let the Netflix client stream video from Netflix servers. In between, we introduce an additional hop, that uses PHY cooperation/SPA, emulated using the Click Modular Router toolkit [16]. The Click toolkit runs as a kernel module on the client node. The kernel module intercepts 802.11 frames arriving from the network interface card, filters out frames from the Netflix servers, and passes them through our emulation link before forwarding them to the higher layers. To other applications on the client, Click is a transparent layer in the network stack.

⁵The bit-rate is limited so that degradation in performance can directly be observed in video streaming. Recommended Netflix bit-rate for HD streaming is 5 Mbps.

PHY Emulation (Testbed). The WARP testbed, described in Section 5.2, is used to send over-the-air (OTA) frames over a network with three relays assisting a source destination pair. SPA is used to adapt to the best cooperation mode to use. The Testbed generates traces describing how each frame was delivered:

- (0) Frame received successfully using direct transmission from source to destination.
- (1) Frame received successfully using selected PHY cooperation mode after failure of the direct transmission.
- (2) Frame not received successfully using direct transmission and PHY cooperation.

The Testbed uses a payload of 7776 bits corresponding to air times of 180 μ secs and 192 μ secs for the direct transmission and PHY cooperation, respectively.

MAC-layer. The network interface card for Alice is set to a maximum transmission unit (MTU) of 780, to be in accordance with our WARP experiments MTU and to avoid additional overhead in TCP that can arise due to dropping fragmented frames. We verified, through TCP flags, that Netflix adjusts to this MTU configuration using Path MTU Discovery. Our MAC layer retransmission policy is similar to 802.11. A transmission is considered in error (requiring retransmission) if the frame was received incorrectly over both two timeslots (direct and cooperation transmissions). In this case, a retransmission is requested (which consequentially incurs more delay). The maximum number of MAC retransmission is 2. After two retries, the frame is dropped.

Click. We use Click to implement the MAC-layer policy described above. The OTA frame-by-frame traces collected from the WARP Testbed are used by Click to introduce delay(s) and drops to Netflix frames, as would have been experienced by the MAC Layer at Alice when using the Testbed for the PHY cooperation hop:

Delay 1 - No Error: If the frame was correctly received in the first timeslot's direct transmission. It incurs a transmission delay equal to a single direct transmission.

Delay 2 - No Error: If the frame was not correctly received in the first timeslot, but was received correctly through cooperation in the second. The delay added is equivalent to direct and PHY cooperation airtimes.

Delay 2 - Error: If the frame was received incorrectly




Figure 12: Video streaming performance of PHY cooperation with SPA.

through both the direct transmission in the first timeslot, and cooperation in the second timeslot.

If a Netflix frame is assigned an error then Click applies the required delays and then views the next frame trace; if again it experiences an error, then that frame is dropped by Click. Otherwise, the corresponding delay is applied and then the frame is forwarded to the higher layer in the device. We make the assumption, which is valid for our indoor scenario, that the propagation delay is negligible in comparison to the transmission time. The emulated link is modeled similar to an ad-hoc 802.11 mode, with disabled RTS and CTS.

Since our delay is emulated, delay jitter is a concern. To meet tight timing delays, we time our delays with spinlocks on a dedicated processor in the node. As a result, in 99.5% of the frames, the jitter was zero (rounded to the nearest μ secs) and only 0.01% of the frames suffered from inaccuracies beyond 5% of the target delay.

Why Emulation? (1) With the objective of comparing our scheme to adaptive routing, we decided to equip adaptive routing with clairvoyance-i.e., genie-aided knowledge of *future* frame drops and retransmission requests, while providing *no clairvoyance* to PHY cooperation/SPA. This way, we provide our competitor with an unfair advantage that we cannot otherwise provide with a real time implementation. Restricting ourselves to a comparison with any *specific* routing protocol would make the case for PHY cooperation/SPA weaker. (2) To faithfully compare the video streaming gains using PHY cooperation/SPA vs adaptive routing, it is necessary to be able to create similar channel conditions for both tests. This is done by time-multiplexing frames that use PHY cooperation and routing in the testbed. A real time experiment would need to run schemes sequentially, and hence it is not feasible to maintain identical channels conditions over long experiments (of the order of tens of seconds). (3) A frame-by-frame OTA trace (from the testbed) of the PHY FERs is all that is needed by the MAC and higher layers, owing to the layered structure of networks. While giving us the advantages mentioned above, we are able to provide a fair comparison across all schemes in this manner.

7.2 Video Streaming Performance

We performed long WARP OTA experiments over multiple 3-relay configurations (running all 6 modes per configuration). These configurations create a time-varying sample of 3-relay networks (with a set number of topology transitions across time) over which we emulate video streaming. Frame-by-frame traces from these OTA experiments are used by click to implement the MAC-layer with delays, retransmissions and drops.

Adaptation and cooperation. Fig. 12(a) plots the bitrates streamed from Netflix, when the Click emulated last hop uses the traces from the 6 fixed modes, as well as SPA. There were 20 different 3-relay topologies in the sample used to generate the frame-by-frame traces, with a total of 132000 frames, each. We see in Fig. 12(a) that SPA is able to maintain the top 2 quality levels (3000 Kbps and 2350 Kbps) most of the time, while each of the 6 fixed modes suffers a significant hit in streaming rates. Moreover, the performance gap between the fixed modes and SPA increases as time progresses, because, once the video rate drops due to a momentarily poor mode connection, it does not instantly get restored when the physical channels recover for two reasons: (1) the slow-start algorithm used by TCP restarts when frames are dropped, and (2) the adaptation strategy in Netflix gradually progresses to higher rates, as opposed to making abrupt jumps. SPA avoids these by offering consistent performance. Fig. 12(b) plots average streaming bitrates, and shows that SPA provides over 95% average video throughput improvement over the best of the six fixed modes (The non-cooperative direct-transmission (DT) scheme is not plotted as it was mostly in “buffer empty” stalling state).

Adaptive PHY beats adaptive routing. One can ask: why not just use an adaptive IP-level routing protocol, that simply finds the best path to connect S and D , and uses the relays for routing as opposed to PHY layer cooperation? To answer this, we put our adaptive PHY cooperation strategy to the test, against a per-packet genie-aided routing protocol that we implement over our Click-emulated link. The genie-aided protocol selects the best possible route, given advance knowledge

Strategy	Packet drop rate
$S - R_1 - D$	27.63 %
$S - R_2 - D$	13.70 %
$S - R_3 - D$	11.44%
genie-aided routing	1.60 %
SPA + PHY Cooperation	0.023 %

Table 2: Packet drop rates for video streaming.

of all future frame traces, and thus outperforms every implementable protocol. In particular:

- Each link allows up to 4 max MAC retransmissions.
 - If a packet can be delivered to D through any path, it selects the path that requires the least total number of retransmissions, and thus deliver the packet faster.
 - If the packet cannot be delivered - given the maximum retransmissions at each link - send the packet along the route that drops it first, and thus, incurs the least delay.
- Table 2 also shows the reduction in packet drop ratio when using SPA over the genie-aided protocol.

Fig. 12(c) plots the bitrates streamed from Netflix (for a 50 sec video play time) with genie-aided routing and SPA. Genie-aided routing makes a very slow start and is at most able to support the lowest video quality; in contrast, SPA allows Netflix to quickly start streaming and gradually improve the quality of video streamed, reaching the highest quality of 3000 kbps towards the end of the play time, achieving an overall $6\times$ throughput gain. This can be explained by the fact that routing along a path requires the relay to perfectly decode; SPA requires only the receiver to do so, leading to lower rates of frame delivery failures and consequently less frame drops. The genie-aided routing protocol suffers from higher frame drops, which causes TCP/IP to misinterpret it as network congestion and therefore degrade its end-to-end throughput.

8. CONCLUSIONS

We presented the design and experimental evaluation of adaptively selecting a small subset of relays to assist an S-D communication via PHY cooperation. Our main conclusions are: (1) The returns from using multiple relays for PHY cooperation diminishes progressively; a well-chosen (small) subset extracts a large portion of the potential gains, while keeping complexity and overhead demands reasonable. (2) For practical PHY cooperation, using network channels to select the highest capacity subnetwork is not a good choice (unlike for routing). (3) We get significant benefits, and a consistent performance across topology changes, by having SPA adaptively select which sub-network to use. (4) PHY cooperation with SPA shows a $6\times$ gain in Netflix streaming throughput, over genie-aided best-path routing (that does not provide PHY cooperation).

APPENDIX

A. OUTAGE PROBABILITY DERIVATION

The approximate capacity of the N relay diamond network with a direct link (within an additive constant) is the minimum over 2^N mutual information terms.

$$\bar{C} \approx \min_{\Omega} I(X_{\Omega}; Y_{\Omega^c} | X_{\Omega^c})$$

Each cut $\Omega \subseteq \{1, 2, \dots, N\}$, and the mutual information for the corresponding cut is denoted by I_{Ω} from here on. As shown in [1, 20], I_{Ω} can be approximated as follows in terms of the capacity of the $2N + 1$ links in the network:

$$I_{\Omega} \approx \max \left\{ \begin{array}{l} \log(1 + h_{sd}^2), \\ \left\{ \max_{i \in \Omega} \{\log(1 + h_i^2)\} + \max_{j \in \Omega^c} \{\log(1 + g_j^2)\} \right\} \end{array} \right\} \quad (1)$$

where h_{sd} , h_i and g_j denote the *absolute values* of the $S - D$, $S - R_i$ and $R_j - D$ channels. The signal and noise powers are normalized to unity.

For a Rayleigh fading channel model, $h_{sd}^2 \sim \mathbf{Exp}(\lambda)$, $h_i^2 \sim \mathbf{Exp}(\lambda_i^l)$ and $g_j^2 \sim \mathbf{Exp}(\lambda_j^r)$. Given a rate R , the probability that \bar{C} is less than R , i.e., the outage probability, can be routinely upper bounded as:

$$P_{\text{out}}(R, h_{sd}, \{h_i, g_i\}_{i \in [1:N]}) = \Pr\left\{\bigcup_{\Omega} I_{\Omega} < R\right\} \leq \sum_{\Omega} \Pr\{I_{\Omega} < R\}$$

In practice, the upper bound is fairly tight. From (1),

$$\Pr\{I_{\Omega} < R\} \stackrel{(a)}{\approx} \frac{\Pr\{\log(1 + h_{sd}^2) < R\}}{\times \Pr\left\{\begin{array}{l} \log(1 + \max_{i \in \Omega} \{h_i^2\}) \\ + \log(1 + \max_{j \in \Omega^c} \{g_j^2\}) \end{array} < R\right\}}$$

where (a) is true because h_{sd} is independent of the h_i 's and g_j 's. Since $h^2 \sim \mathbf{Exp}(\lambda)$, we have

$$\Pr\{\log(1 + h_{sd}^2) < R\} = 1 - e^{-\lambda(2^R - 1)}$$

For the second term in the product, $X \triangleq \max_{i \in \Omega} \{h_i^2\}$ and $Y \triangleq \max_{j \in \Omega^c} \{g_j^2\}$. Then the second term becomes

$$P_{\Omega} = \Pr\{\log(1 + X) + \log(1 + Y) < R\}$$

If $F_X(x)$ is c.d.f of X and $f_X(x)$ is its p.d.f,

$$P_{\Omega} = \int_{x=0}^{2^R-1} f_X(x) F_Y\left(\frac{2^R}{1+x} - 1\right) dx$$

Now, $F_X(x)$ and $f_X(x)$ (similarly $F_Y(y)$, $f_Y(y)$) are:

$$F_X(x) \stackrel{(b)}{=} \prod_{i \in \Omega} \Pr\{h_i^2 \leq x\} = \prod_{i \in \Omega} (1 - e^{-\lambda_i^l x})$$

$$f_X(x) = \frac{d}{dx} F_X(x) = \sum_{i \in \Omega} \lambda_i^l e^{-\lambda_i^l x} \prod_{k \in \Omega \setminus \{i\}} (1 - e^{-\lambda_k^l x})$$

where (b) is true since h_i 's are independent. P_{Ω} can now be written as a (weighted) sum of a constant and several

integrals as $P_{\Omega}^t = \int_{x=0}^{2^R-1} e^{-\alpha x} e^{-\beta(\frac{2^R}{1+x}-1)} dx$. Then,

$$P_{\text{out}}(R, h_{sd}, \{h_i, g_i\}_{i \in [1:N]}) \approx (1 - e^{-\lambda(2^R-1)}) \sum_{\Omega} P_{\Omega}$$

For numerical evaluation, we essentially need to *compute* terms of the type P_{Ω}^t above.

Optimization problem: Given N relays, we want to select an *outage optimal* subnetwork with k relays ($k < N$). To do so, we need to solve the following problem:

$$P_{\text{out}}^{\text{opt}}(R, k) = \arg \min_{\mathcal{I}: |\mathcal{I}|=k} P_{\text{out}}(R, h_{sd}, \{h_i, g_i\}_{i \in \mathcal{I}}) \quad (2)$$

B. REFERENCES

- [1] AVESTIMEHR, A. S., DIGGAVI, S. N., TIAN, C., AND TSE, D. N. An approximation approach to network information theory. *Foundations and Trends in Communications and Information Theory* 12, 1-2 (2015), 1–183.
- [2] AVESTIMEHR, A. S., DIGGAVI, S. N., AND TSE, D. N. C. Wireless network information flow: A deterministic approach. *IEEE Trans. on Inf. Theory* 57, 4 (2011), 1872–1905.
- [3] BALAN, H. V., ROGALIN, R., MICHALOLIAKOS, A., PSOUNIS, K., AND CAIRE, G. AirSync: Enabling distributed multiuser MIMO with full spatial multiplexing. *IEEE/ACM Trans. on Netw.*, 99 (2013).
- [4] BEJARANO, O., AND KNIGHTLY, E. Virtual MISO triggers in Wi-Fi-like networks. In *INFOCOM* (2013), pp. 1582–1590.
- [5] BRAHMA, S., DUARTE, M., SENGUPTA, A., WANG, I.-H., FRAGOULI, C., AND DIGGAVI, S. QUILT: A decode/quantize-interleave-transmit approach to cooperative relaying. In *INFOCOM* (2014), pp. 2508–2516.
- [6] CHANG, K. ET AL. Relay operation in 802.11ad. In *IEEE 802.11ad TGad 1-/0494r1* (2010).
- [7] DEB, S., MHATRE, V., AND RAMAIYAN, V. WiMAX relay networks: Opportunistic scheduling to exploit multiuser diversity and frequency selectivity. In *MOBICOM* (2008), pp. 163–174.
- [8] DUARTE, M., SENGUPTA, A., BRAHMA, S., FRAGOULI, C., AND DIGGAVI, S. Quantize-map-forward (QMF) relaying: An experimental study. In *MOBIHOC* (2013), pp. 227–236.
- [9] GOLLAKOTA, S., PERLI, S. D., AND KATABI, D. Interference alignment and cancellation. In *ACM SIGCOMM Computer Communication Review* (2009), vol. 39, ACM, pp. 159–170.
- [10] GUAN, Z., MELODIA, T., YUAN, D., AND PADOS, D. Distributed spectrum management and relay selection in interference-limited cooperative wireless networks. In *MOBICOM* (2011), pp. 229–240.
- [11] HERBSTER, M., AND WARMUTH, M. K. Tracking the best expert. *Machine Learning* 32, 2 (1998), 151–178.
- [12] JAIN, S., AND DAS, S. R. Exploiting path diversity in the link layer in wireless ad hoc networks. *Ad Hoc Networks* 6, 5 (2008), 805–825.
- [13] JING, Y., AND H., J. Single and multiple relay selection schemes and their achievable diversity orders. *IEEE Transactions on Wireless Communications* 8, 3 (March 2009), 1414–1423.
- [14] KATTI, S., GOLLAKOTA, S., AND KATABI, D. Embracing wireless interference: Analog network coding. In *SIGCOMM* (2007), pp. 397–408.
- [15] KHATTAB, A., CAMP, J., HUNTER, C., MURPHY, P., SABHARWAL, A., AND KNIGHTLY, E. W. WARP: A flexible platform for clean-slate wireless medium access protocol design. *SIGMOBILE Mob. Comput. Commun. Rev.* 12, 1 (2008), 56–58.
- [16] KOHLER, E., MORRIS, R., CHEN, B., JANNOTTI, J., AND KAASHOEK, M. F. The click modular router. *ACM Transactions on Computer Systems* 18, 3 (2000), 263–297.
- [17] LO, C. K., VISHWANATH, S., AND HEATH, R. W. Relay subset selection in wireless networks using partial decode-and-forward transmission. *IEEE Transactions on Vehicular Technology* 58, 2 (Feb 2009), 692–704.
- [18] MURPHY, P., AND SABHARWAL, A. Design, implementation, and characterization of a cooperative communications system. *IEEE Trans. on Vehic. Tech.* 60, 6 (2011), 2534–2544.
- [19] NAZAROGLU, C., OZGUR, A., AND FRAGOULI, C. Wireless network simplification: The gaussian n-relay diamond network. In *ISIT* (2011), pp. 2472–2476.
- [20] PAWAR, S., AVESTIMEHR, A., AND TSE, D. Diversity-multiplexing tradeoff of the half-duplex relay channel. In *Allerton* (2008), pp. 27–33.
- [21] RAHUL, H., HASSANIEH, H., AND KATABI, D. SourceSync: a distributed wireless architecture for exploiting sender diversity. In *SIGCOMM* (2010), pp. 171–182.
- [22] SHI, Y., SHARMA, S., HOU, Y., AND KOMPELLA, S. Optimal relay assignment for cooperative communications. In *MOBIHOC* (2008), pp. 3–12.
- [23] WANG, B., HAN, Z., AND LIU, K. Distributed relay selection and power control for multiuser cooperative communication networks using stackelberg game. *IEEE Trans. on Mob. Comp.* 8, 7 (2009), 975–990.
- [24] YUKSEL, M., AND ERKIP, E. Broadcast strategies for the fading relay channel. In *Military Communications Conference, 2004. MILCOM*

2004. *2004 IEEE* (2004), vol. 2, IEEE, pp. 1060–1065.